An Equilibrium Model with Computationally Constrained Agents[1]

Wolfgang Kuhle

Max Planck Institute for Research on Collective Goods, Kurt-Schumacher-Str. 10, 53113 Bonn, Germany. Email: kuhle@coll.mpg.de.

arXiv:1611.01771v1 [q-fin.EC] 6 Nov 2016

Abstract: *We study a large economy in which firms cannot compute exact solutions to the non-linear equations that characterize the equilibrium price at which they can sell future output. Instead, firms use polynomial expansions to approximate prices. The precision with which they can compute prices is endogenous and depends on the overall level of supply. At the same time, firms' individual supplies, and thus aggregate supply, depend on the precision with which they approximate prices. This interrelation between supply and price forecast induces multiple equilibria, with inefficiently low output, in economies that otherwise have a unique, efficient equilibrium. Moreover, exogenous parameter changes, which would increase output were there no computational frictions, can diminish agents' ability to approximate future prices, and reduce output. Our model therefore accommodates the intuition that interventions, such as unprecedented quantitative easing, can put agents into "uncharted territory".*

Keywords: Polynomial Inference, Self-Referential Equilibria, Glitch Equilibria

1 Introduction

Few people would claim that they are able to compute future equilibrium outcomes, such as prices, with any accuracy. Despite this, textbook models implicitly assume that, given all relevant data, agents compute exact numeric values for future equilibrium prices, respectively, the entire distribution of these prices if the model involves risk. In this paper, we assume that economic agents are computationally constrained to the use of polynomial functions. That is, instead of being able to solve arbitrary non-linear problems, they rely

[1]I am particularly indebted to Martin Hellwig for detailed comments on an earlier draft of this paper. I also thank Dominik Grafenhofer, Sebastian Klein, Harvey Lapan, and Carl Christian von Weizsäcker for discussions on equilibrium models. Finally, I received helpful questions and comments from seminar participants in Bonn. First draft July 2015.

on polynomial expansions to approximate future equilibrium outcomes. Put differently, agents act just like economic researchers who use polynomials, such as the Arrow-Pratt approximation, to restate complicated non-linear problems in terms of workable polynomials.

Using a two-period model, in which firms employ polynomial approximations to infer future selling prices for their output, we find that multiple equilibria emerge in well-behaved economies that would have a unique, efficient equilibrium if agents could compute future equilibria with perfect accuracy. Moreover, exogenous parameter changes, which would increase economic activity were agents computationally unconstrained, can reduce economic activity as they make it harder for agents to approximate equilibrium prices.

Our results rely on the fact that there are levels of supply where a polynomial approximation to the function, which relates equilibrium supply to equilibrium price, is of good quality, and other levels where it is of low quality. Put differently, a firm's ability to compute equilibrium prices changes with the level of aggregate supply. At the same time, individual supply, and thus aggregate supply, varies with the precision with which agents can predict prices. This interaction gives rise to two coexisting types of equilibria. In the first, economic activity falls into intervals where agents' polynomial approximations are of high quality and the role of the computational friction is small. These equilibria can coincide with the rational expectations equilibrium (REE). In the second type, computational frictions are important and agents find it difficult to predict prices: Aggregate supply is (i) low and (ii) falls into an interval where agents' approximations, to the equation describing equilibrium, are of poor quality.

In one interpretation, we may think of a farmer who must decide in spring how much corn he should plant. This farmer may know the price at which corn tends to sell in years with "normal" supply. Moreover, he might know that small increases in aggregate supply tend to reduce prices, i.e, that demand is locally downward sloping. Finally, he may know that this downward slope tapers off as supply increases. The farmer, however, is unable to calculate all numeric values that the demand function takes over its entire domain. If he wants to calculate those prices, which obtain once supply differs from those levels that he is familiar with, he must use a polynomial expansion to extrapolate the new price. In a macroeconomic interpretation, we think of a large number of firms that have to choose production today in anticipation of future demand. These firms know the price at which their goods sell in "normal times". However, if firms collectively cut production today, it will be difficult for them to know whether future selling prices

increase, due to reduced supply, or fall, since the layoffs, associated with production cuts, reduce demand.[2] This intuition extends naturally to economies where demand concerns a vector of goods, which may involve substitutes and complements. In such a setting it appears even more natural to assume that firms cannot solve for the overall equilibrium. Instead, a firm, which produces a particular good, may, if it is exceptionally well informed, use the economy's Jacobian matrix to compute demand for its particular good in terms of a first-order polynomial approximation.

Regarding parameter changes, the uncertainty that agents face in our model does not originate from a world with stochastically changing parameters. Instead, agents know the magnitude of the parameter change in advance; the difficulty is to predict its consequences. As an illustration, we refer to two representative comments made on the quantitative easing program: Stanley Druckenmiller, an accomplished investor with a thirty-year track record, commented in retrospect "I didn't know how it was going to end... I would have said inflation [which] would have been dead wrong."[3] Similarly, taking an ex-ante perspective, Joseph Stiglitz predicted that QE2 would likely bring interest rates down "a little", but that it was unclear what the risks, ranging from economic growth to"a whole set of other potential risks that - may result from this policy"[4], were.[5] Likewise, there appears to be no consensus among observers how, if at all, a UK exit from the EU will affect the economies of the UK and the remainder EU. Similar arguments apply to more long-term problems such as the demographic transition.[6] As we show, agents' inability to perform comparative statics alters model predictions considerably: Large parameter changes, which would unambiguously increase output if agents could compute the model's

[2]Diamond (1982), and Cooper and John (1988), develop models where search frictions result in upward-sloping aggregate demand functions. See Hellwig (1993) for a review of models with non-monotonic demand. More generally, such effects are important in economies where Say's law, stating that supply creates its own demand, is of relevance. See also Chamley (2014) for a related model of savings and investment.

[3]Speech given at the 2015 Dealbook conference.

[4]Interview given on the Charlie Rose show on 3 November 2010.

[5]See Soros (1994) for case studies highlighting market participants' difficulties to anticipate the impact that pre-announced central bank policies have on macroeconomic equilibrium variables. Similarly, Niederhoffer (1997), p. 381, concludes his discussion on excess demand functions: "The difficulty is that nobody knows what the equilibrium level is until at least the morning after the fact."

[6]That is, we know from birth statistics that cohorts entering the labor market will be smaller and cohorts entering retirement will grow. At the same time, it proves difficult to predict how such changes impact future growth paths.

comparative statics correctly, can reduce output. The model therefore accommodates the intuition that interventions, such as unprecedented quantitative easing, can put agents into "uncharted territory", i.e., diminish agent's ability to forecast relevant equilibrium variables. Hence, even though markets work inefficiently, the government's ability to improve market outcomes is limited.

Related Literature: Due to their bounded computational capacity, agents work with an approximate, misspecified model. Except for special cases, they cannot form rational expectations in the sense of Hutchison (1937), Grunberg and Modigliani (1954), Muth (1961), Blanchard (1979) and DeCanio (1979), which are consistent with the true model. Regarding model misspecification, our approach is thus akin to the literature on learning, Bray (1982), Marcet and Sargent (1989), and Sargent (1993), where agents use a misspecified least squares approach to infer unknown model parameters. Rothschild (1974) and McLennan (1984) model firms that experiment with different supply functions to learn about stochastic demand. Firms in our model are small, and thus changing individual supply has no influence on prices. That is, in the dynamic extension of our model, the information that firms learn over time is determined by overall equilibrium rather than individual experimentation.

We interpret our baseline model as a simple A^D, A^S setting, as in Keynes (1936) and Samuelson (2009), which is augmented with a computational friction. In Section 4, we show that our model may be reinterpreted as a Diamond (1982) and Cooper and John (1988) aggregate search model, where the probability of finding a trading partner depends on the equilibrium level of economic activity. In this interpretation, agents' computational constraint makes it difficult for them to compute the equilibrium probability of finding a trading partner. Regarding government intervention, Diamond (1982) and Cooper and John (1988) find positive multipliers, which are due to the search friction. In the current model, where the search friction is coupled with a computational friction, government intervention has a non-monotonous effect on output.

Tesfatsion (2006), Gintis (2007), Farmer and Foley (2009), and Thurner et al. (2012) argue for "agent-based" models in which agents follow decision rules that do not necessarily coincide with rational behavior.[7] In the current paper, agents are computationally

[7]One argument for such a departure is computational complexity: Rubinstein (1998), Maymin (2011) and Ackerman et al. (2011) emphasize that agents might be constraint in their ability to count or to compute conditional probabilities. Nelson and Winter (1985), Ulanowicz (2008), Hofbauer and Sandholm (2011), Arthur (2015), and Kuhle (2016) for evolutionary models where biases emerge endogenously.

constrained to the use of polynomial expansions. This case is of special interest since researchers in economics, physics, and engineering indeed rely on first- and second-order polynomial expansions, rather than exact solutions, to understand non-linear problems; equilibrium comparative statics of a well behaved model, $y = f(y; b)$, are commonly evaluated in terms of a first-order polynomial expansion $\Delta y \approx \Delta y f_y + \Delta b f_b$, which yields $\frac{\Delta y}{\Delta b} \approx \frac{f_b}{1-f_y}$, respectively, $\frac{\partial y}{\partial b} = \frac{f_b}{1-f_y}$.[8]Hence, we argue that the current model is methodologically consistent in the sense that the outside researcher, i.e., the paper's reader, will use the same method of analysis that is used by the model's agents.

Section 2 abstracts from computational frictions and identifies the unique rational expectations equilibrium. Section 2.1 studies equilibria that obtain with computationally constrained agents. Section 2.2 considers the impact of exogenous parameter changes. Section 3.1 studies the economy's convergence to the rational expectations equilibrium in a dynamic setting, where firms accumulate empirical knowledge. Section 3.2 introduces asymmetric information. In Section 4, we suggest different interpretations of our baseline model. Section 5 concludes.

2 Model

We study a large economy in which a mass one of firms $i \in [0, 1]$ produce a homogenous good. There are two periods of time. In the first period, each firm chooses to produce a quantity of goods a_i in anticipation of a future selling price $\hat{P}$. In the second period, firms sell the finished products a_i inelastically to consumers at a market clearing price P. For simplicity, to ensure uniqueness of the REE in the economy without computational friction, we assume that aggregate demand is twice continuously differentiable and

[8]Similarly, Mas-Colell et al. (1995), pp. 599-641, use first-order polynomial expansions to examine non-linear demand in pure exchange economies with many commodities. As mentioned earlier, Finetti (1952), Arrow (1970), Pratt (1964) use second-order expansions to study expected utility. Likewise, the familiar first- and second-order conditions, $f'(x_0) = 0$ and $f''(x_0) < 0$, for a smooth function $f(x)$ to have local maximum at point x_0, stem from an expansion $f(x) = f(x_0)+f'(x_0)(x-x_0)+\frac{1}{2}f''(x_0)(x-x_0)^2+O^3$; see Chiang and Wainwright (2005), pp. 250-253, or Samuelson (1947), pp. 357-379. Related, the stability of differential and difference equations, Samuelson (1947), pp. 21-121, 257-349, and 380-439, or Galor (2007). Finally, DeCanio (1979), p. 52, points out that solving expectations equilibria requires either that the model is assumed to be linear, or that the model's equations have to be linearized, i.e., rewritten in terms of a first-order polynomial, which is what our agents do.

monotonously downward-sloping in goods quantity A:

$$P = \phi(A), \quad \phi_A < 0, \quad \phi_{AA} \gtrless 0, \quad \phi(0) > 0. \tag{1}$$

Demand (1) represents the model's non-linearity, respectively, the computational obstacle that agents have to overcome. In Section 4, we suggest three different interpretations of $\phi()$ by showing that it captures the non-linearities that individual agents face when they make forward-looking decisions in the standard workhorse models of Diamond (1965), Diamond (1982), and the A^D, A^S model of Samuelson (2009). That is, ϕ may be interpreted as (i) future returns to savings, (ii) the probability of finding a trading partner, or (iii) the selling price for output.

Firm i chooses a production schedule a_i^* to maximize expected profits

$$a_i^* = \arg\max_{a_i} \Big\{ \pi_i = a_i \hat{P} - \frac{1}{2} a_i^2 \Big\}, \quad a_i \geq 0,$$

where $\hat{P}$ is the firm's expectation regarding the selling price and $\frac{1}{2}a_i^2$ is a quadratic cost function. Hence, agent i supplies

$$a_i^* = \hat{P}$$

and aggregate supply is

$$A = \int_{[0,1]} a_i^* di = \hat{P}. \tag{2}$$

If agents are computationally unconstrained, they can compute demand (1) over its entire domain. That is, for each level of aggregate supply A, they form rational price expectations $\hat{P} = \phi(A)$. In turn, they combine (1) and (2) to calculate the unique equilibrium quantity A_0:

$$A_0 = \phi(A_0), \tag{3}$$

and, using (1), they compute equilibrium price P_0:

$$P_0 = \phi(A_0). \tag{4}$$

Accordingly, we have

Lemma 1. *There exists a unique, rational expectations equilibrium* $\{A_0, P_0\} \in \mathcal{R}^2_+$. *In this equilibrium agents forecast prices correctly* $\hat{P} = P_0$.

Proof. Market clearing (3)-(4) determines equilibrium quantity $A_0 > 0$, which is unique since $\phi(0) > 0$ and $\phi_A < 0$. Using (1), the equilibrium price is $P_0 = \phi(A_0) = A_0 > 0$. Finally, (2) indicates that $\hat{P} = A_0$ and thus $\hat{P} = P_0$. □

2.1 Polynomial Equilibria

We now assume that firms cannot compute demand over its entire domain. Instead, they are familiar with a point on the demand function, $A^*, \phi(A^*)$, and the demand function's slope $\phi_A(A^*)$ at this point. It is convenient to start with the assumption that this point is the REE of Lemma 1, i.e, agents know $A_0, \phi(A_0)$, and the slope $\phi_A(A_0)$. In turn, once supply differs from A_0, agents use polynomial expansions to extrapolate demand to estimate the resulting price. Polynomial equilibrium points will be those points where the polynomial, which mimics true demand, intersects with supply. Put differently, "polynomial equilibria" are those points that solve the agents' approximate model. The REE from the previous section will be one, but not the only, such equilibrium.

Expanding demand (1) around the perfect foresight equilibrium, agents forecast the selling price (1) as:

$$\hat{P} = \phi(A_0) + \phi_A(A_0)\Delta A, \quad \Delta A = A - A_0. \tag{5}$$

Equation (5) reflects that agents cannot numerically compute the true price $P = \phi(A_0 + \Delta A)$ at which a supply $A = A_0 + \Delta A$ sells. The reliability of estimate (5) decreases the more aggregate supply A differs from A_0.[9] We assume that agents choose output a_i to maximize:

$$a_i^* = \arg\max_{a_i}\Big\{\pi = a_i\hat{P} - \frac{1}{2}a_i^2 - a_i\tau\Delta A^2\Big\}, \quad \tau \geq 0. \tag{6}$$

The profit criterion (6) allows for two interpretations. In the first, $a_i\hat{P} - \frac{1}{2}a_i^2$ is the firm's profit given the price estimate $\hat{P}$, and $-a_i\tau\Delta A^2$ reflects that firms, knowing their estimate is based on a first-order expansion, which neglects second-order terms, discount $\tau > 0$ the estimated revenue. In a second interpretation, which we elaborate on in Proposition 2 of Appendix A, $-\tau$ represents the demand function's second derivative ϕ_{AA}. In this case agents do not discount their price estimate, and rely on a second-order Taylor-series expansions to estimate the selling price.[10]

[9] Note that the model's coefficients throughout can be chosen such that the equilibrium deviation ΔA is arbitrarily small, respectively, such that the approximation (5) is of arbitrarily good quality.

[10] That is, if agents knew demand's second derivative, their price estimate (5) would write $\hat{P} = \phi(A_0) + \phi_A\Delta A + \frac{1}{2}\phi_{AA}\Delta A^2$. Substituting this into (6), and setting the discount rate $\tau = 0$, yields $a_i^* = \arg\max_{a_i}\Big\{a_i(\phi(A_0) + \phi_A\Delta A + \frac{1}{2}\phi_{AA}\Delta A^2) - \frac{1}{2}a_i^2\Big\}$. Comparison indicates that the new profit criterion is equivalent to the old, (6), except for the second-order derivative $\frac{1}{2}\phi_{AA}$ taking the place of the discount rate $-\tau$.

From (6), we obtain individual and aggregate supply:

$$a_i^* = \hat{P} - \tau \Delta A^2, \quad A = \int_{i \in [0,1]} a_i di = \hat{P} - \tau \Delta A^2. \tag{7}$$

Combining supply (7) and estimated demand (5), the equilibrium quantity A, where supply intersects with the demand estimate, is the solution to:

$$A = \phi(A_0) + \phi_A(A_0)\Delta A - \tau \Delta A^2. \tag{8}$$

For convenience, we identify equilibria $j = 0, 1, 2...$, in terms of their distance $\Delta A_j = A_j - A_0$ to the rational expectations equilibrium A_0. That is, $\Delta A_j = 0$ corresponds to the REE. Using the fact that $A_0 = \phi(A_0)$, we rewrite (8) as:

$$\tau \Delta A^2 + (1 - \phi_A)\Delta A = 0,$$

and note:

Proposition 1. *There exists the rational expectations equilibrium* $\Delta A_0 = A_0 - A_0 = 0$ *in which agents' price forecasts are correct* $\hat{P} = P_0$. *There exists a second equilibrium* $\Delta A_1 = A_1 - A_0 = -\frac{(1-\phi_A)}{\tau} < 0$ *in which* $\hat{P}_1 \gtreqless P_1$.

Both equilibria in Proposition 1 are self-fulfilling. In the perfect foresight equilibrium, no firm deviates from the equilibrium supply $A = A_0$, and thus there is no need for agents to rely on polynomial approximations: Producers *know* the price $\phi(A_0)$. The opposite is the case in the second equilibrium: Once firms supply $A_1 \neq A_0$, they are uncertain as to the equilibrium price, $\phi(A_1) = \phi(A_0 + \Delta A_1)$, which they can only approximate as $\phi(A_0) + \phi_A(A_0)\Delta A$. Moreover, the error of this approximation, $(A - A_0)^2$, grows the more agents deviate from supplying A_0. That is, once firms deviate from the rational expectations equilibrium, they find it harder to estimate future prices and thus they are incentivised to deviate even further until a new equilibrium is reached. In this equilibrium, firms cannot forecast prices accurately, and thus they choose to produce a small number of goods, at a low marginal cost, which provides a margin of safety.

As we argued earlier, this interdependence between aggregate output and the individual firm's ability to understand the environment that it operates in is a crucial aspect in most crises: Once consumers and investors change their behavior, they find themselves in an environment that is hard to understand, and they hold back on investment and consumption decisions waiting for the "dust to settle". In the current interpretation, by cutting production, agents put *themselves* into "uncharted territory". This aspect

is, by assumption, not captured in environments where agents can compute the entire demand function, respectively, solve the model as in Lemma 1. Before we discuss the scope for government to correct such "glitches" in output, which turns out to be limited, we make one remark: The model's coefficients $\phi_A(A_0), \tau$ can be chosen such that $\Delta A_1 = A_1 - A_0 = -\frac{(1-\phi_A)}{\tau} < 0$ is arbitrarily small. That is, both equilibria in Proposition 1 exist even if the the first-order Taylor-series approximation (5) is of very good quality, i.e., if the error term is of order $\mathcal{O}(\Delta A^2)$.

2.2 Parameter Changes

We augment demand $P = \phi(A; b)$ to incorporate an exogenous parameter b. This parameter is assumed to increase demand $\phi_b = \phi_b(A; b) > 0$ for every A. This parameter may be seen as government demand or money supply.[11] In this interpretation, the following section identifies the multiplier effect that obtains once agents need to rely on approximations to anticipate the consequences of policy interventions.

We begin with a benchmark model where agents are computationally unconstrained. Second, we study the model with friction. Comparing both settings shows that parameter increases, which increase demand and equilibrium output in a model with unconstrained agents, can reduce economic activity if firms are computationally constrained. Put differently, parameter changes, in particular if they are large, can put agents into "uncharted territory", and incentivise them to cut, rather than increase, output.

2.2.1 Comparative statics without friction

Recalling our augmented demand function:

$$P = \phi(A; b_0), \quad \phi_A < 0, \quad \phi_{AA} \gtrless 0, \quad \phi_b > 0, \quad \phi(0; b) > 0, \tag{9}$$

firms can anticipate the equilibrium price P correctly, if they are computationally unconstrained as in Lemma 1. Hence, they choose a production schedule a_i^* which maximizes profits $a_i^* = \arg\max_{a_i} \left\{ \pi_i = a_i P - \frac{1}{2} a_i^2 \right\}$. Aggregate supply is thus

$$A = \int_{[0,1]} a_i^* di = P. \tag{10}$$

[11] Alternatively, as we discuss in Section 4, the model may be interpreted as the capital market of an overlapping generations economy, where a_i, A, $\phi()$ are, respectively, individual savings and aggregate savings, and $\phi(A)$ is the marginal product of capital that agents *expect* to receive on their savings. Finally, b may be seen as public debt and A_0 as steady state capital.

Taken together (9) and (10) yield a unique equilibrium P_0, A_0 for *every* given exogenous parameter b_0. Once the parameter changes from b_0 to $b_1 = b_0 + \Delta b$, the price is again correctly anticipated as the unique solution P_1, A_1 to the equations $P = A$ and $P = \phi(A; b_0 + \Delta b)$.

How would an actual human being, or an economic researcher, try to think about the impact of the parameter change? The outside researcher, who uses textbook methods to study how changes in the exogenous parameter from b_0 to b_1 change output and price, cannot compute A and P explicitly. Instead, he will approximate the model's comparative statics. That is, he will differentiate (9) and (10):

$$\Delta P \approx \phi_A(A_0; b_0)\Delta A + \phi_b(A_0; b_0)\Delta b, \quad \Delta A = A - A_0, \quad \Delta b = b_1 - b_0, \tag{11}$$

$$\Delta A = \Delta P. \tag{12}$$

Combining (11) and (12) yields the model's comparative statics:

Lemma 2. *Exogenous parameter variations Δb change output (and price) according to $\frac{\Delta A}{\Delta b} \approx \frac{\phi_b}{1-\phi_A} > 0$ and $\frac{\partial A}{\partial b} = \lim_{\Delta b \to 0} \frac{\Delta A}{\Delta b} = \frac{\phi_b}{1-\phi_A} > 0$.*

That is, an outside observer/analyst would use a polynomial expansion of $A = \phi(A; b), P = \phi(A; b)$ to *approximate* the impact of an exogenous parameter change as in Lemma 2. In the following section, we assume that firms themselves make such "polynomial inference" using such an approximation to anticipate the consequences of parameter changes.

2.2.2 Comparative statics with computationally constrained agents

Using a first-order expansion, around $A_0 = \phi(A_0; b_0)$, agents approximate the equilibrium price:

$$\hat{P} = \phi_0(A_0, b_0) + \phi_A \Delta A + \phi_b \Delta b, \quad \Delta A = A - A_0, \quad \Delta b = b_1 - b_0. \tag{13}$$

That is, agents have to incorporate two aspects in their demand forecast: (i) the direct effect of the parameter change Δb and (ii) the equilibrium response of all agents who deviate $\Delta A \neq 0$ from their usual supply choice. As before, agents discount the price estimate since they do not know how curvature terms of demand ϕ_{AA}, ϕ_{bb}, and ϕ_{Ab} affect

prices:[12]

$$\pi_i = a_i\hat{P} - a_i(\tau_1\Delta A^2 + \tau_2\Delta b^2 + \tau_3|\Delta A||\Delta b|) - \frac{1}{2}a_i^2$$
$$a_i^* = \hat{P} - \tau_1\Delta A^2 - \tau_2\Delta b^2 - \tau_3|\Delta A||\Delta b| \quad \tau_i \geq 0, \quad i = 1, 2, 3. \tag{14}$$

To find the equilibria associated with (13) and (14), it is useful to distinguish cases where $\Delta A \geq 0$ from cases where $\Delta A \leq 0$.

We begin by looking for equilibria where $\Delta A \geq 0, \Delta b \geq 0$. If $\Delta A \geq 0$ and $\Delta b \geq 0$ then supply equals approximate demand (14), when:

$$A = \phi_0(A_0, b_0) + \phi_A\Delta A + \phi_b\Delta b - \tau_1\Delta A^2 - \tau_2\Delta b^2 - \tau_3\Delta A\Delta b,$$

and thus:

$$\Delta A_{1,2} = -\frac{1 - \phi_A + \tau_3\Delta b}{2\tau_1} \pm \sqrt{(\phi_b - \tau_2\Delta b)\frac{1}{\tau_1}\Delta b + \left(\frac{1 - \phi_A + \tau_3\Delta b}{2\tau_1}\right)^2}. \tag{15}$$

Combining the two equilibrium candidates in (15) with our initial assumption $\Delta A \geq 0$, we have:

Lemma 3. *If and only if $\Delta b < \frac{\phi_b}{\tau_2}$, there exists an equilibrium in which, compared to the perfect foresight equilibrium, production (and price) are increased:* $\Delta A_1 = -\frac{1-\phi_A+\tau_3\Delta b}{2\tau_1} + \sqrt{(\phi_b - \tau_2\Delta b)\frac{1}{\tau_1}\Delta b + \left(\frac{1-\phi_A+\tau_3\Delta b}{2\tau_1}\right)^2} > 0$. *At the margin, increases in the parameter increase income if* $\frac{\partial\Delta A}{\partial\Delta b} = -\frac{\tau_3}{2\tau_1} + 1/2\frac{\phi_b\frac{1}{\tau_1} - 2\frac{\tau_2}{\tau_1}\Delta b + 2\frac{\tau_3}{\tau_1}\left(\frac{1-\phi_A+\tau_3\Delta b}{2\tau_1}\right)}{\sqrt{(\phi_b - \tau_2\Delta b)\frac{1}{\tau_1}\Delta b + \left(\frac{1-\phi_A+\tau_3\Delta b}{2\tau_1}\right)^2}} > 0.$

Proof. Follows directly from (15). □

To interpret the equilibrium in Lemma 3, we study how it corresponds to the REE of Proposition 1. That is, we note that $\lim_{\Delta b \to 0} \Delta A_1(\Delta b) = 0$, i.e., the equilibrium quantity A_1 converges to the REE quantity A_0 as $b \to b_0$. On the contrary, for large changes $\Delta b > 0$, the equilibrium loses its RE character as agents do not precisely know how demand is impacted by the exogenous change. Such large parameter changes have an ambiguous effect on output, which is captured by the term $(\phi_b - \tau_2\Delta b)\Delta b$. On the one hand, agents extrapolate the increase in demand $\phi_b\Delta b > 0$. At the same time, agents cannot rule out that too large an increase might eventually prove counterproductive

[12] In Appendix B we discuss an alternative error term $\max[\Delta A^2, \Delta b^2, \Delta A\Delta b]$, where agents are either concerned about miscalculating the parameter change's impact, $\Delta b > \Delta A$, or the other agents' reaction $\Delta A > \Delta b$ to the parameter change.

$-\tau_2 \Delta b^2$. That is, large, unprecedented changes in the model's structure render agents' polynomial approximations unreliable, and put them into "uncharted territory".

Lemma 3 thus features four policy regimes. First, if the parameter change is (infinitesimally) small, the agents' polynomial approximations are of high quality. In this regime, policy is as effective as in the model of Section 2.2.1, Lemma 2, where agents are computationally unconstrained. That is, the model's multiplier is given by $\frac{\partial \Delta A}{\partial \Delta b}_{|\Delta b=0} = \frac{\partial A}{\partial b} = \frac{\phi_b}{1-\phi_A} > 0$. Second, there is an intermediate region where $\Delta b \in [0, \Delta b_1]$, and policy changes have a positive effect at the margin, $\frac{\partial \Delta A}{\partial \Delta b} > 0$. These marginal returns, however, are diminishing. Third, there is a region $\Delta b \in [\Delta b_1, \Delta b_2]$ where agents find themselves in "uncharted territory" and start to cut production $\frac{\partial \Delta A}{\partial \Delta b} < 0$. Finally, in the extreme case where $\Delta b \geq \frac{\phi_b}{\tau_2}$, an equilibrium $\Delta A > 0$ cannot exist.

Regarding the remaining equilibria, where $\Delta A \leq 0$, we recall (13) and (14), and note that $-|\Delta A| = \Delta A$. Accordingly, there are two candidates

$$\Delta A_{2,3} = -\frac{1-\phi_A-\tau_3\Delta b}{2\tau_1} \pm \sqrt{(\phi_b - \tau_2 \Delta b)\frac{1}{\tau_1}\Delta b + \Big(\frac{1-\phi_A-\tau_3\Delta b}{2\tau_1}\Big)^2}. \tag{16}$$

In view of (16), if $\Delta b < \frac{\phi_b}{\tau_2}$, there exists exactly one equilibrium, in which $\Delta A_2 = -\frac{1-\phi_A-\tau_3\Delta b}{2\tau_1} - \sqrt{(\phi_b - \tau_2 \Delta b)\frac{1}{\tau_1}\Delta b + \Big(\frac{1-\phi_A-\tau_3\Delta b}{2\tau_1}\Big)^2} < 0$ such that economic activity is lower than in the REE. For large parameter changes, $\Delta b > \frac{\phi_b}{\tau_2}$, there can exist up to two equilibria in which economic activity is low. Combining these observations with Lemma 3 yields:

Corollary 1. *Large parameter changes $\Delta b > \frac{\phi_b}{\tau_2}$ preclude the existence of equilibria with output $A \geq A_0$.*

3 Extensions

So far agents were assumed to know the rational expectations equilibrium $A_0, \phi(A_0)$. In Section 3.1, we study a dynamic setting where agents learn different points on the demand curve over time. In turn, we examine how the economy converges to the REE. Second, in our baseline setting, all agents know the same point on the demand curve. Price forecasts and supply decisions are therefore the same across agents. Once different agents know different pieces of the demand curve, this is no longer true. Rather than knowing each other's price forecasts and supplies, agents have to estimate price *and* supply

simultaneously. In Section 3.2, we extend our model to incorporate such asymmetric information in a manner which is akin to Bayesian inference.

3.1 Learning

We abstracted from the fact that agents may learn from past mistakes, i.e., suboptimal production choices that were based on incorrect price estimates. One would imagine that they memorize these mistakes, or the observation that a quantity A_1 is associated with an observable price $P_1 = \phi(A_1)$. Under our current assumptions on the demand function, this price differs from the estimated price $\hat{P}_1$. Hence, agents would not supply A_1 again. Second, if agents are computationally constrained to the use of polynomials, how do they find the perfect foresight equilibrium in the first place? This section's main observation is that agents will learn the REE over time.

Agents who sell repeatedly into the market will, over time $t = 0, 1, 2, 3...$, observe an increasing number of points on the demand curve. Regarding these points, $P_t = \phi(A_t)$, we assume that agents also learn demand's slope $\phi_A(A_t)$ once a quantity A_t is marketed. Given past observations $A_t, t = 0, 1, 2, 3, ..., T$ agents can refine their price estimate as:

$$\hat{P}_{T+1} = \phi(A^*) + \phi_A(A^*)(A_{T+1} - A^*), \quad A^* = \arg\min_{A_t}\Big\{|A_{T+1} - A_t|\Big\} \quad t = 0, 1, 2, 3...T \tag{17}$$

That is, to estimate prices, they select from the set of known points $\{A_t\}_{t=0}^T$ the point A^*, which is closest to the future supply A_{T+1}. Put differently, they use the observation A^* from the past, which is most similar/closest to the situation they are trying to make inference on. In turn, agents i choose supply

$$a_i = \hat{P}_{T+1} - (A_{T+1} - A^*)^2. \tag{18}$$

Hence, for a given A^*, there are two equilibrium candidates

$$A_{T+1} = A^* - \frac{1 - \phi_A(A^*)}{2} \pm \sqrt{(\phi(A^*) - A^*) + \Big(\frac{1 - \phi_A(A^*)}{2}\Big)^2}. \tag{19}$$

To show that (17) and (19) ensure that agents learn the REE equilibrium $A_0, \phi(A_0)$, we proceed in two steps. First, we study the case where demand $\phi()$ is a convex function. In this case, convergence to the REE can be studied in terms of a simple first-order difference equation. Second, for the remaining cases, we give an indirect argument in Appendix C.

3.1.1 Convex Demand

Without loss of generality, we assume that agents start with a prior $\mu, \phi(\mu)$, $\mu < A_0$. Moreover, we focus on the "+" roots of (19). For convex demand, we now show that (17) and (19) imply a first-order difference equation for supply:

$$A_{T+1} = A_T - \frac{1-\phi_A(A_T)}{2} + \sqrt{(\phi(A_T) - A_T) + \Big(\frac{1-\phi_A(A_T)}{2}\Big)^2}. \tag{20}$$

First, we note that (20) indicates that agents, in Marshallian fashion, increase supply, such that $A_{T+1} > A_T$, if the marginal revenue $\phi(A_T)$ exceeds the marginal cost of production A_T. Moreover, from Lemma 1, we know that there exists only one level of supply, namely A_0, where $A = \phi(A)$. It follows that (20) has a unique steady state at the point where $A_{T+1} = A_T = A_0$. This steady state equilibrium is stable due to our assumption that ϕ is downward-sloping: For $A_T < A_0$, we have $\phi(A_T) - A_T > 0$ and thus $A_{T+1} > A_T$. At A_0, the system is locally stable since $\frac{\partial A_T}{\partial A_{T+1}} =_{|A_T = A_0} 0 \in (-1, 1)$. To complete the argument, we note that the sequence $\{A_t\}_{t=1}^T$ is strictly increasing, and, due to our convexity assumption on ϕ, that $A_t \leq A_0 \forall t = 0, 1, 2....T$.[13] That is, A_T always adjusts towards, but not beyond, A_0. Hence, according to (17), agents will always use the information that they learned in the previous period, when a quantity A_T was marketed, to think about A_{T+1}. This last property allows us to study convergence in terms of the first-order difference equation (20).

3.2 Asymmetric Information

One may suspect that heterogeneity in information might mitigate the possibility of multiple equilibria that we emphasize.[14] Moreover, one might expect that dispersion of private information induces some agents to supply too little and others too much such that, on average, errors cancel and supply might actually be at an efficient level.[15] Regarding

[13] To see this, recall (17) and (18), which imply $A_{T+1} = \phi(A_T) + \phi_A(A_T)(A_{T+1} - A_T) - (A_{T+1} - A_T)^2 \leq \phi(A_T) + \phi_A(A_T)(A_{T+1} - A_T)$. At the same time, convexity of ϕ implies: $\phi(A_{T+1}) = \phi(A_T + A_{T+1} - A_T) \geq \phi(A_T) + \phi_A(A_T)(A_{T+1} - A_T)$. Taking both inequalities together, we have $A_{T+1} - \phi(A_{T+1}) \leq 0$, respectively, $A_{T+1} \leq A_0$. Where $A_{T+1} \leq A_0$, follows from ϕ being downward-sloping and $A = \phi(A)$ at $A = A_0$. Hence, if we start at a point $\mu - \phi(\mu) < 0$, this implies that $A_{T+1} \leq A_0 \forall T$. Finally, as mentioned before, if demand is non-convex, A_t can overshoot A_0. In that case we require an additional argument, which we give in Appendix C.

[14] See, e.g., Morris and Shin (1998) for a coordination problem where asymmetric information selects unique equilibria in an economy with a continuum of players.

[15] See Galton (1907), and Grossman and Stiglitz (1976) for such wisdom of the crowd effects.

these conjectures, we find that (i) multiplicity carries over to the case with dispersed private information and (ii) that dispersed information tends to amplify (dampen) supply if demand is concave (convex).

Each agent $i \in [0,1]$ is assumed to know the selling price $\phi(A_i)$ and demand's first derivative $\phi_A(A_i)$ of a particular supply $A_i \in [0,\infty]$. Agents are distributed over these points according to an integrable density function $f()$. For simplicity, we normalize agents' discount rate to $\tau = 1$.

Conditional on information $A_i, \phi(A_i), \phi_A(A_i)$ agent i supplies:

$$a_i^*|A_i = \hat{P}|A_i - (\hat{A}|A_i - A_i)^2, \tag{21}$$

where $\hat{P}|A_i$ and $\hat{A}|A_i$ are agent $i's$ price and supply forecasts conditional on knowing demand $\phi(A_i)$ at point A_i. The polynomial estimates for price and quantity are:

$$\hat{P}|A_i = \phi(A_i) + \phi_A(A_i)(\hat{A}|A_i - A_i), \tag{22}$$

and

$$\hat{A}|A_i = \int_{[0,1]} (a_j|A_j)|A_i dj. \tag{23}$$

Where (23) reflects that agent i uses his information at point A_i to infer the information and thus the supply of the other agents who know a different point A_j. That is, *agent i knows that agent j observes a point on the same demand curve and thus he uses a polynomial expansion around $\phi(A_i)$ to estimate the information $\phi(A_j), \phi_A(A_j)$ that player j receives*. Based on this reasoning, i can construct an estimate for the other players' price estimates, which he needs to calculate aggregate supply. Agent $i's$ price and supply estimates are thus given by the simultaneous solution of (22)-(23). In turn, he can choose supply (21). We solve the model in Appendix E using a guess-and-verify approach.

These solutions yield two main insights. First, as in Proposition 1, multiple equilibria exist due to the interaction between agents' ability to forecast the equilibrium and aggregate supply. Second, unlike the earlier model, where information was symmetric, we show in Lemma 4 that output is depressed across *all* equilibria since agents systematically underestimate demand if the true demand function is convex. Moreover, the marginal cost of production differs among producers, and thus output, in addition to being low, is produced inefficiently.

4 Interpretations

In this section, we reinterpret our model in terms of the three macroeconomic workhorse frameworks: (i) aggregate search models of the Diamond (1982) type, (ii) Life-cycle savings models of the Diamond (1965) type, and (iii) models of supply and demand as in Samuelson (2009) and Mas-Colell et al. (1995).

1) Search: In the context of the Diamond (1982), p. 887, model, agents face the following choice problem:

$$\max_{a_i}\{a_i\phi(A;b_0)-f(a_i)\},\quad a_i\in\mathcal{R}_+.$$

Where a_i is individual $i's$ output choice in Period 1, and $\phi(A;b_0)$ is the probability with which agents find a trading partner in Period 2. If a trading partner is found, agent i can sell/exchange goods at price one. Finally, $f(a_i)$ is the cost of production.

The chance of finding a trading partner, $\phi(A;b_0)$, is an increasing function in aggregate economic activity $A=\int_{[0,1]}a_i di$ and government demand b. Diamond (1982) shows that such an economy can have multiple REE equilibria, which we call, say, A_0, A_1. Suppose now that agents know one of these REE, e.g., A_0; then, if they are computationally constrained, as in the present paper, they would need to use a polynomial expansion to compute the probability $\phi(A_0+\Delta A)$ of finding a trading partner that would prevail once agents collectively deviate ΔA from producing A_0. The same applies to the evaluation of the exogenous policy parameter b, which may, unlike in Diamond (1982), result in a negative multiplier effect, as discussed in Section 2.2.

2) Savings and Investment: In the context of the Diamond (1965) model, $\phi()$ may be interpreted as a component to agents' consumption savings problem:

$$\max_{s_t,c_t^1,c_{t+1}^2}U(c_t^1,c_{t+1}^2)\quad s.t.\quad c_t^1=w_t-s_t,\quad c_{t+1}^2=s_t(1+r_{t+1}),\quad c_t^1>0,c_{t+1}^2>0,(24)$$

where factor prices are functions of the prevailing capital-labor ratios k_t and k_{t+1}:

$$r_{t+1}=f'(k_{t+1}),\quad w_t=f(k_t)-f'(k_t)k_t.$$

To make choice (24), agents have to form expectations regarding equilibrium interest $r_{t+1}=f'(k_{t+1})$. In equilibrium, the life-cycle savings condition, $(1+n)k_{t+1}=s_t$, relates savings and capital; n representing the exogenous rate of population growth.

Suppose now that the economy is initially in a steady state at k_0, s_0. To compute future interest rates, agents have to compute $r_{t+1} = f'(k_{t+1}) = f'(\frac{s_t}{1+n}) =: \phi(s_t; n)$. Once again, if agents know the prevailing interest in the steady state, they have to engage in polynomial expansions to form price expectations $\hat{r}_{t+1}$ to compute the interest rate that obtains in the (temporary) equilibria that obtain once agents choose savings $s_t = s_0 + \Delta s_t$.

This argument extends to the case where agents supply labor in both periods. In that case, to make their savings decision, agents have to (i) approximate the interest rate and (ii) the second-period wage rate. That is, they have to approximate the Samuelson (1962) neoclassical factor-price frontier, $w_{t+1} = \xi(r_{t+1}(k_{t+1}(s_t)))$, which relates wages to interest, interest to the capital intensity, and finally the capital intensity to savings.

3) Supply and Demand: Our lead interpretation was that of a simplified A^D, A^S setting. Taking this perspective, we suggest one reason why demand analysis may be computationally complicated for firms. Suppose demand is given by $A^D = \xi(A; L(A))$, where ξ represents demand, which is downward-sloping in supply $\frac{\partial \xi}{\partial A} < 0$. At the same time, cuts in production A reduce employment $L(A)$ and demand, i.e., $\frac{\partial \xi}{\partial L}\frac{\partial L}{\partial A} > 0$. Accordingly, agents grapple with the question whether the function $\phi(A) := \xi(A; L(A))$ is upward or downward sloping once supply falls into regions which they do not know from past experience.

Similarly, suppose that demand concerns a vector $\mathbf{A} \in \mathcal{R}^L$ of goods, involving substitutes and complements, as in Mas-Colell et al. (1995), pp. 599-641. Suppose that a firm produces a particular good a_l, then, if all other firms in the economy, supplying the various other goods, change their behavior, it has to analyze how changes in the supplies and prices for the other goods $L \setminus l$, influence demand, and thus price, for good l. In turn, if this firm is exceptionally well-informed, it might know the entries of the economy's Jacobian matrix. However, it need not know demand's second- and cross-derivatives, which once again makes it difficult to compute demand correctly.

5 Conclusion

Economists' forecasting record suggests that it is difficult to compute future economic events. The current model recognizes this and assumes that agents cannot compute exact numeric values for future equilibrium outcomes such as prices.

The model's key feature is that the precision with which agents can approximate fu-

ture equilibrium prices depends on the level of aggregate economic activity, and is thus endogenous. This interdependence between aggregate output and an individual firm's ability to forecast the price at which it can sell its output gives rise to equilibria in which economic activity is inefficiently low. Such equilibria may be interpreted as "glitches" of the overall economy. During such a glitch, agents collectively reduce economic activity. This change in behavior makes it difficult to forecast the resulting equilibrium, which, in turn, justifies the initial output cut. For similar reasons we also find that the scope for government to correct such "glitches" in output is limited: Interventions, which would unambiguously increase output in a frictionless economy, can make it harder for firms to predict future equilibria and reduce output even further. Our model therefore captures the common place observation that large parameter changes, such as unprecedented quantitative easing, can put agents into "uncharted territory".

The particular form in which equilibria obtain depends on the assumption that agents use the same Taylor-series expansions that an economic researcher, who applies standard textbook methods, would use. More sophistication on the part of agents will undoubtedly change the specific form and number of equilibria. However, it appears unlikely that the precision with which future equilibrium outcomes are approximated can ever be entirely independent of the overall level of economic activity, which is what our findings rely on.

A Second-order Expansions

In this appendix, we derive our results for a setting where agents know of demand's first and second derivatives. They can thus use second-order expansions to estimate prices:

$$\hat{P} = \phi(A_0) + \phi_A \Delta A + \frac{1}{2}\phi_{AA}\Delta A^2, \tag{25}$$

We study the equilibria that emerge once agents are averse $\tau > 0$ to the third-order error. Setting $\tau = 0$, we obtain the equilibria that emerge if agents are indifferent regarding errors. Using the estimate (25), firms choose a profit-maximizing quantity:

$$a_i^* = \arg\max_{a_i} \Big\{ \pi = a_i \hat{P} - \frac{1}{2}a_i^2 - a_i \tau |\Delta A^3| \Big\}, \quad \tau \geq 0.$$

Individual and aggregate supply are thus

$$a_i^* = \hat{P} - \tau |\Delta A|^3, \quad A = \int_{[0,1]} a_i^* di = \hat{P} - \tau |\Delta A|^3. \tag{26}$$

Combining (25) and (26), we obtain:

$$A + \tau |\Delta A|^3 = \phi(A_0) + \phi_A \Delta A + \frac{1}{2}\phi_{AA}\Delta A^2. \tag{27}$$

Recalling that $A_0 = \phi(A_0)$, (27) writes:

$$\tau |\Delta A|^3 - \frac{1}{2}\phi_{AA}\Delta A^2 + (1 - \phi_A)\Delta A = 0. \tag{28}$$

If $\tau = 0$ we have:

Proposition 2. *If $\tau = 0$, there exists the perfect foresight equilibrium $\Delta A_0 = 0$ and a second equilibrium $\Delta A_1 = \frac{1-\phi_A}{\frac{1}{2}\phi_{AA}}$.*

According to Proposition 2, economic activity in the polynomial equilibrium exceeds activity in the perfect foresight equilibrium if the demand function's second derivative indicates that demand may rebound $\frac{1}{2}\phi_{AA} > 0$ once supply exceeds A_0. A negative second derivative $\frac{1}{2}\phi_{AA}$ depresses output in the same manner as the discount factor τ in Proposition 1. As before, estimated demand can meet supply more than twice if agents discount their price estimate:

Proposition 3. *If $\tau > 0$, there exist at least two equilibria: $\Delta A_0 = 0$, and $\Delta A_1 = -\frac{1}{4\tau}\phi_{AA} - \sqrt{(\frac{1}{4\tau}\phi_{AA})^2 + \frac{1-\phi_A}{\tau}} < 0$, in which economic activity is depressed. If $\phi_{AA} > 0$, and $(\frac{1}{4\tau}\phi_{AA})^2 > \frac{1-\phi_A}{\tau}$, there may exist two additional equilibria where economic activity is elevated $\Delta A_{2,3} = \frac{1}{4\tau}\phi_{AA} \pm \sqrt{(\frac{1}{4\tau}\phi_{AA})^2 - \frac{1-\phi_A}{\tau}} > 0$.*

Proof. Using (28), we find the perfect foresight equilibrium $\Delta A_0 = 0$. To identify the remaining equilibria, we distinguish cases (i) $\Delta A < 0$ and (ii) $\Delta A > 0$.

1.) Assuming $\Delta A < 0$: we note that $|\Delta A^3| = -\Delta A^3$. Dividing by ΔA, we find that (28) has roots: $\Delta A_{1,2} = -\frac{1}{4\tau}\phi_{AA} \pm \sqrt{(\frac{1}{4\tau}\phi_{AA})^2 + \frac{1-\phi_A}{\tau}}$. However, only one root satisfies the initial assumption $\Delta A < 0$. That is, $\Delta A = -\frac{1}{4\tau}\phi_{AA} - \sqrt{(\frac{1}{4\tau}\phi_{AA})^2 + \frac{1-\phi_A}{\tau}} < 0$, regardless of the sign of ϕ_{AA}. The other root $\Delta A = -\frac{1}{4\tau}\phi_{AA} + \sqrt{(\frac{1}{4\tau}\phi_{AA})^2 + \frac{1-\phi_A}{\tau}} > 0$ is positive since $(\frac{1}{4\tau}\phi_{AA})^2 > 0$, $\phi_A < 0$, $\frac{1-\phi_A}{\tau} > 0$. It thus violates the initial assumption $\Delta A < 0$.

2.) Assuming $\Delta A > 0$: we note that $|\Delta A^3| = \Delta A^3$ and find that (28) has two real roots $\Delta A_{1,2} = \frac{1}{4\tau}\phi_{AA} \pm \sqrt{(\frac{1}{4\tau}\phi_{AA})^2 - \frac{1-\phi_A}{\tau}}$ if $(\frac{1}{4\tau}\phi_{AA})^2 > \frac{1-\phi_A}{\tau}$. Both of these roots are negative if $\phi_{AA} < 0$ violating the assumption $\Delta A > 0$. Hence, $\phi_{AA} > 0$ is a sufficient condition for $\Delta A > 0$ equilibria to exist. □

The polynomial equilibria in propositions 1-4 have in common that they originate from a coordination problem: In their price forecasts, each agent takes the overall supply A as given. In equilibrium, however, aggregate supply depends on agents' price forecasts. Hence, the price forecast itself is an equilibrium outcome. At this point, it is clear that higher-order polynomials yield even more equilibria, and that propositions 1-4 carry over qualitatively once we introduce demand $\Phi(\mathbf{A})$, for a vector $\mathbf{A}$ of goods.

B Alternative Discounting

The price estimate is as before:

$$\hat{P} = \phi_0(A_0, b_0) + \phi_A \Delta A + \phi_b \Delta b, \quad \Delta A = A - A_0, \quad \Delta b = b - b_0. \tag{29}$$

The model differs in agents' discounting:

$$\begin{aligned} \pi_i &= a_i \hat{P} - a_i \max[\Delta A^2, \Delta b^2, \Delta A \Delta b] - \frac{1}{2} a_i^2 \\ a_i^* &= \hat{P} - \max[\Delta A^2, \Delta b^2] \end{aligned} \tag{30}$$

According to (30) there are two regimes. First, agents fear that they miscalculate the impact of the exogenous parameter change in case $\Delta b > \Delta A$. Second, agents fear that they misjudge the other agents' reaction to the exogenous parameter variation $\Delta A > \Delta b$. To analyze the equilibrium outcomes associated with (29) and (30), we distinguish cases where the parameter change is relatively large, $|\Delta b| > |\Delta A|$, from cases where its impact is relatively small, $|\Delta b| < |\Delta A|$. Note that $|\Delta A|$ is a function of $|\Delta b|$. That is, we start

with the assumption that, e.g., $|\Delta b| > |\Delta A|$ and solve for the equilibrium ΔA. In turn, we check whether the initial assumption $|\Delta b| > |\Delta A|$ is correct.

Proposition 4. *If* $|\frac{\phi_b - \Delta b}{1-\phi_A}| < 1$*, there exists only one equilibrium* $\Delta A = \left(\frac{\phi_b}{1-\phi_A} - \frac{\Delta b}{1-\phi_A}\right)\Delta b$ *in which* $|\Delta b| > |\Delta A|$*. In this equilibrium* $\lim_{b \to b_0} \frac{\Delta A}{\Delta b} = \frac{\phi_b}{1-\phi_A} = \frac{\partial A}{\partial b}_{|A_0, b_0}$.

Proof. Individual supply is $a_i = \hat{P} - \max[\Delta A^2, \Delta b^2]$ under the assumption that $|\Delta b| > |\Delta A|$, we have $a_i = \hat{P} - \Delta b^2$. Aggregate supply is thus $A = \int_{i \in [0,1]} a_i = \hat{P} - \Delta b^2$. Equilibrium requires $A = A_0 + \phi_A \Delta A + \phi_b \Delta b - \Delta b^2$, respectively, $\Delta A(1 - \phi_A) + \Delta b(\Delta b - \phi_b) = 0$. Solving yields $\Delta A = \frac{\phi_b - \Delta b}{1-\phi_A} \Delta b$. It remains to note that $|\frac{\phi_b - \Delta b}{1-\phi_A}| < 1$ ensures that $|\Delta b| > |\Delta A|$. □

As before, Proposition 4, $\Delta A = \left(\frac{\phi_b}{1-\phi_A} - \frac{\Delta b}{1-\phi_A}\right)\Delta b$, shows that agents extrapolate the positive impact that a parameter change $\Delta A = \frac{\phi_b}{1-\phi_A}$. At the same time, they are facing increased uncertainty as to the actual price at which their products sell $\Delta A = -\frac{\Delta b}{1-\phi_A}\Delta b$. A large parameter change where $\Delta b > \phi_b$ thus reduces economic activity as the uncertainty that it creates outweighs the expansive effect $\phi_b > 0$.

This leaves us with equilibria where agents are more concerned about the potential error associated with aggregate supply changes. In these cases, agents are primarily afraid that they forecast prices incorrectly due to the change ΔA. For cases where $|\Delta b| < |\Delta A|$ we have:

Proposition 5. *There exists an upper bound* $\Delta b_1 > 0$ *and an equilibrium where* $\Delta A_1 = -\frac{1}{2}(1 - \phi_A) - \sqrt{\phi_b \Delta b + (\frac{1}{2}(1 - \phi_A))^2} < 0$ *and* $|\Delta b| < |\Delta A|$*, if* $\Delta b \in [0, \Delta b_1]$*. If* $\frac{\phi_b}{1-\phi_A} > 1$ *there exists an upper bound* $\Delta b_2 > 0$ *and a second equilibrium* $\Delta A_2 = -\frac{1}{2}(1 - \phi_A) + \sqrt{\phi_b \Delta b + (\frac{1}{2}(1 - \phi_A))^2} > 0$ *where* $|\Delta b| < |\Delta A|$*, if* $\Delta b \in [0, \Delta b_2]$.

Proof. Individual supply is $a_i = \hat{P} - \max[\Delta A^2, \Delta b^2]$ under the assumption that $|\Delta b| < |\Delta A|$, we have $a_i = \hat{P} - \Delta A^2$. Aggregate supply is thus $A^S = \int_{i \in [0,1]} a_i = \hat{P} - \Delta A^2$. Equilibrium requires $A^S = A^D$ such that $A = A_0 + \phi_A \Delta A + \phi_b \Delta b - \Delta A^2$, respectively, $\Delta A^2 + (1 - \phi_A)\Delta A + \phi_b \Delta b = 0$. Solving yields $\Delta A_{1,2} = -\frac{1}{2}(1 - \phi_A) \pm \sqrt{\phi_b \Delta b + (\frac{1}{2}(1 - \phi_A))^2}$. Both roots are real since we assumed $\phi_b > 0$ and $\Delta b > 0$. It remains to specify the conditions under which our initial hypothesis $|\Delta b| < |\Delta A|$ holds. We start with $\Delta A_1 = -\frac{1}{2}(1 - \phi_A) - \sqrt{\phi_b \Delta b + (\frac{1}{2}(1 - \phi_A))^2}$ and note (i) $\Delta A_1(\Delta b = 0) = -(1 - \phi_A)$ such that $|\Delta b| = 0 < |\Delta A|$, (ii) the derivative $\frac{\partial \Delta A}{\partial \Delta b} = \frac{-\frac{\phi_b}{2}}{\sqrt{\phi_b \Delta b + (\frac{1}{2}(1-\phi_A))^2}}$ vanishes as Δb becomes large. Taken together, (i) and (ii) imply that an upper bound Δb_1 exists, such that

$|\Delta b| < |\Delta A|$ as long as $\Delta b \in [0, \Delta b_1]$. Similarly, regarding the second equilibrium $\Delta A_2 = -\frac{1}{2}(1-\phi_A) + \sqrt{\phi_b \Delta b + (\frac{1}{2}(1-\phi_A))^2}$, we note that (i) $\Delta A_2(\Delta b = 0) = 0$ such that $|\Delta b| = |\Delta A| = 0$, (ii) the derivative $\frac{\partial \Delta A}{\partial \Delta b} = \frac{\frac{\phi_b}{2}}{\sqrt{\phi_b \Delta b + (\frac{1}{2}(1-\phi_A))^2}} =_{|\Delta b=0} \frac{\phi_b}{(1-\phi_A)}$ and $\lim_{\Delta b \to \infty} \frac{\partial \Delta A}{\partial \Delta b} = \lim_{\Delta b \to \infty} \frac{\frac{\phi_b}{2}}{\sqrt{\phi_b \Delta b + (\frac{1}{2}(1-\phi_A))^2}} = 0$. Taken together, (i) and (ii) imply that if $\frac{\phi_b}{1-\phi_A} > 1$ there exists an upper bound Δb_2 such that $|\Delta b| < |\Delta A|$ as long as $\Delta b \in [0, \Delta b_2]$. □

The first equilibrium in Proposition 5 corresponds to the perfect foresight equilibrium, $\Delta A = 0$ of Proposition 1: In the limit, where the parameter change becomes infinitesimally small, we obtain $\Delta A = 0$. In this equilibrium, increases in b indeed increase equilibrium supply provided that these increases are small such that $\Delta b < \Delta b_1 < \Delta b_2$. In the second equilibrium, which corresponds to the crisis equilibrium in Proposition 1, output is strictly decreasing in b.

Taken together, Propositions 4 and 5 suggest that bold interventions by the government tend to reduce economic activity as such changes make it more difficult for agents to forecast prices. Moreover, in the crisis equilibrium ΔA_2, government interventions, which would increase output were there no computational frictions, always reduce income.

C Learning The REE

We have shown that agents can learn the A_0 equilibrium if (i) demand is convex and (ii) agents always coordinate on the "+" equilibrium. We now show that agents also learn the REE if (i) demand is not convex, and (ii) when agents, e.g., alternate between playing the "+" and "-" root equilibria.

We write the equilibrium condition as:

$$A - \phi(A^*) = \phi_A(A^*)(A - A^*) - (A - A^*)^2. \tag{31}$$

The left-hand side of (31) represents the difference between supply and demand in an equilibrium where agents use their knowledge of $A^*, \phi(A^*)$ to estimate the price. This difference is 0 in the perfect foresight equilibrium where $A_0 = \phi(A_0)$. Regarding the right-hand side, we define $\varepsilon = (A - A^*)$, which yields

$$A - \phi(A^*) = \phi_A \varepsilon - \varepsilon^2, \tag{32}$$

over time as agents observe an increasing number of price quantity pairs $\{\phi(A_t), A_t\}_{t=0}^T$. The distance $\varepsilon = (A - A^*)$ between the aggregate supply A and the point of estimation

A^* will go to zero.[16] Hence, from (32) we have $\lim_{\varepsilon \to 0}(A - \phi(A^*)) = 0$. However, the only point where $A = \phi(A)$ is A_0, the perfect foresight equilibrium quantity. Hence, as agents learn more data on demand, they eventually move towards the efficient equilibrium.

D Asymmetric Equilibria

Agents using $\phi(A^*)$ to think about a deviation from A^*, might choose a quantity $A(A^*)$ which is closer to A^{**} than to A^*. Thus, to forecast the selling price $\phi(A(A^*))$, they would rather use A^{**} as the point of departure. In turn, once agents use A^{**} they might choose a supply $A(A^{**})$, which, however, is again closer to A^* than A^{**}. Thus, given (17), agents would switch back to A^*, and so forth:

$$|A(A^*) - A^*| > |A(A^*) - A^{**}| \tag{33}$$

$$|A(A^{**}) - A^{**}| > |A(A^{**}) - A^*|. \tag{34}$$

In such a situation, there exists no symmetric equilibrium between points A^* and A^{**}. Instead, there exists an asymmetric equilibrium in which a mass $\psi \in (0, 1)$ of agents use A^* and a mass $1 - \psi$ use A^{**}.

It follows from (17) that agents are only indifferent between using A^* and A^{**} if the equilibrium quantity A satisfies $A = \frac{A^* + A^{**}}{2}$. That is, agents see A^* and A^{**} as equally informative once the supply they want to learn about is equally distant from both points. To establish the existence of an equilibrium we have to prove that there exist shares ψ and $1-\psi$ for which the equilibrium supply $A_\psi \psi + A_{1-\psi}(1-\psi)$ indeed equals $\bar{A} = \frac{A^* + A^{**}}{2}$.

[16] To see this note that, given our assumptions on $\phi()$, all real-valued equilibrium supplies, (20), fall into a compact interval $[0, \hat{A}]$. In turn, the sequence of equilibrium supplies $\{A_t\}_{t=0}^T$, for which agents know demand, partitions this interval. As time progresses, this partition becomes finer and finer. That is, if we order the quantities A_t such that $A_l < A_{l+1}, l = 1, 2, 3...T$, the interval between A_l and A_{l+1} is either filled with new equilibria, or, if there exist no equilibria between them, economic activity takes place elsewhere on the interval $[0, \hat{A}]$. Finally, we note that agents using A^* to think about a deviation from A^*, might choose a quantity $A(A^*)$ which is closer to A^{**} than to A^*. Thus, to forecast the selling price $\phi(A(A^*))$, agents would rather use A^{**} as the point of approximation. In turn, once agents use A^{**} they might choose a supply $A(A^{**})$, which, however, is again closer to A^* than A^{**}. Thus agents would switch back to A^*, and so on. In such situations, there exists no symmetric equilibrium between points A^* and A^{**}. Instead, Appendix D shows that there does exist an asymmetric equilibrium, where a fraction ψ of the agents uses point A^* and the remaining fraction $1 - \psi$ uses point A^{**}, resulting in equilibrium supply $A = \frac{A^* + A^{**}}{2}$.

Once we denote the equilibrium supply by $A(\psi) = \psi A(A^*, A(\psi)) + (1-\psi)A(A^{**}, A(\psi))$, we can write the equations that determine equilibrium as

$$A(\psi) = \bar{A} \quad \bar{A} := \frac{A^* + A^{**}}{2} \tag{35}$$

$$A(\psi) := \psi A_\psi(A^*, A(\psi)) + (1-\psi)A_{1-\psi}(A^{**}, A(\psi)) \tag{36}$$

$$A_\psi(A^*, A(\psi)) = \hat{P}_\psi - (A(\psi) - A^*) \quad A_\psi(A^{**}, A(\psi)) = \hat{P}_{1-\psi} - (A(\psi) - A^{**}) \tag{37}$$

$$\hat{P}_\psi = \phi(A^*) + \phi_A(A^*)(A(\psi) - A^*) \tag{38}$$

$$\hat{P}_{1-\psi} = \phi(A^{**}) + \phi_A(A^{**})(A(\psi) - A^{**}) \tag{39}$$

Solving (36)-(39) for supply $A(\psi)$ yields:

$$A(\psi) = -\frac{p}{2} \pm \sqrt{-q + (\frac{p}{2})^2}$$

$$-q = \psi\phi(A^*) + (1-\psi)\phi(A^{**}) - \psi\phi_A(A^*) - (1-\psi)\phi_A A^{**} - \psi(A^*)^2 - (1-\psi)(A^{**})^2$$

$$p = 1 + 2\psi A^* + 2(1-\psi)A^{**} - \psi\phi_A(A^*) - (1-\psi)\phi_A(A^{**}) > 0$$

Given our assumption that supply is downward-sloping, we have $-p < 0$. Hence, there is only one real root $A(\psi) = -\frac{p}{2} + \sqrt{-q + (\frac{p}{2})^2}$ with positive supply. Note in particular that $A(\psi = 0)$ ($A(\psi = 1)$) is the supply when all agents use A^* (A^{**}) as the point of reasoning. By our initial hypothesis (33)-(34), we have $A(\psi = 0) > \bar{A}$ and $A(\psi = 1) < \bar{A}$. Hence, there exists an intermediate value $\tilde{\psi}$, at which $A(\tilde{\psi}) = \bar{A}$, if $\sqrt{-q + (\frac{p}{2})^2}$ as required by the indifference condition (35).[17]

E Asymmetric Information Equilibria

We solve the model in two steps. First, we guess the equilibrium outcome. Given this guess, we solve the estimation problem for agent i. Second, we solve for the equilibrium and verify our guess.

4) Problem of an individual agent: To find the optimal supply of an individual agent, we must solve (22)-(23). To do so, we start with the *guess* that agent i beliefs that agents j hold the same believe over aggregate supply that he holds himself, i.e., $\hat{A}|A_i = (\hat{A}|A_j)|A_i = \hat{A}$. Given this guess, we show that:

$$\hat{P}|A_i = (\hat{P}|A_j)|A_i. \tag{40}$$

[17]The root $\sqrt{-q + (\frac{p}{2})^2}$ is never complex as ψ runs from 0 to 1. To see this we recall $-q(\psi = 0) > 0$, $-q(\psi = 1) > 0$ and that the derivative $\frac{\partial -q}{\partial \psi}$ is not a function of ψ itself and thus cannot change signs as ψ varies. Accordingly, $-q > 0 \forall \psi \in [0, 1]$. Hence, $\sqrt{-q + (\frac{p}{2})^2}$ is always real.

Put differently, equation (40) means that agent i believes that agents j observe points, which lie on his estimated demand curve. That is, agent i knows that agent j estimates the price as

$$\hat{P}|A_j = \phi(A_j) + \phi_A(A_j)(\hat{A}|A_j - A_j),$$

and thus i estimates $\hat{P}|A_j$ as:

$$(\hat{P}|A_j)|A_i = \phi(A_j)|A_i + \phi_A(A_j)|A_i((\hat{A}|A_j)|A_i - A_j), \tag{41}$$

where

$$\phi(A_j)|A_i = \phi(A_i) + \phi_A(A_i)(A_j - A_i), \tag{42}$$

$$\phi_A(A_j)|A_i = \phi_A(A_i), \quad \hat{A}|A_i = (\hat{A}|A_j)|A_i = \hat{A}. \tag{43}$$

Taken together, (41)-(43) mean that agent i uses polynomials to approximate the demand curve upon which the other player observes a point $A_j, \phi(A_j), \phi_A(A_j)$. Using (42)-(43), (41) rewrites:

$$\begin{aligned}(\hat{P}|A_j)|A_i &= \phi(A_i) + \phi_A(A_i)(A_j - A_i) + \phi_A(A_i)(\hat{A} - A_j),\\ &= \phi(A_i) + \phi_A(A_i)(\hat{A}|A_i - A_i) =_{|(22)} \hat{P}|A_i.\end{aligned}$$

Agent i thus believes that agent j will work with a price estimate that is identical to the one he uses himself and thus he will conclude that $j's$ supply forecast is identical to his own, such that $\hat{A}|A_i = (\hat{A}|A_j)|A_i = \hat{A}$, which confirms our initial guess. It remains to solve (22)-(23) for player $i's$ forecast in the two equilibria $A_{1,2}$:

$$\hat{A}_{1,2}|A_i = A_i - \frac{1-\phi_A(A_i)}{2} \pm \sqrt{(\phi(A_i) - A_i) + \Big(\frac{1-\phi_A(A_i)}{2}\Big)^2},$$

$$\begin{aligned}a_i &= \hat{P}|A_i - (\hat{A}|A_i - A_i)^2\\ &= \phi(A_i) + \phi_A(A_i)(\hat{A}|A_i - A_i) - (\hat{A}|A_i - A_i)^2 \qquad (44)\\ a_{i;1,2} &= \phi(A_i) + \phi_A(A_i)\Big(-\frac{1-\phi_A(A_i)}{2} \pm \sqrt{(\phi(A_i) - A_i) + \Big(\frac{1-\phi_A(A_i)}{2}\Big)^2}\Big)\\ &\quad -\Big(-\frac{1-\phi_A(A_i)}{2} \pm \sqrt{(\phi(A_i) - A_i) + \Big(\frac{1-\phi_A(A_i)}{2}\Big)^2}\Big)^2. \qquad (45)\end{aligned}$$

5) Equilibrium: From (45) we calculate equilibrium supply as:

$$A_{1,2} = \int_{[0,1]} a_{1,2}(A_i) f(A_i) di. \tag{46}$$

Concerning the equilibrium quantity $A_k, k = 1, 2$, we note

Lemma 4. *If demand ϕ is quasi-convex, then equilibrium output across both equilibria $A_k, k = 1, 2$ falls short of efficient output A_0.*

Proof. To compare equilibrium output (46) to efficient output A_0, we recall that $A_0 = \phi(A_0)$. We also recall that, for convex functions f, we have $f'(u) \leq \frac{f(v)-f(u)}{v-u}$. In the current context, this means that agents tend to underestimate convex demand functions:

$$\phi(A_i) + \phi_A(A_i)(A - A_i) \leq \phi(A)$$

recalling (44) we have:

$$a_i = \phi(A_i) + \phi_A(A_i)(\hat{A}|A_i - A_i) - (\hat{A}|A_i - A_i)^2 \leq \phi(A_0) - (\hat{A}|A_i - A_i)^2 \leq \phi(A_0) = A_0.$$

Put differently, agents supply less than the efficient quantity since (i) they underestimate demand and (ii) since they know that their price estimate is inaccurate. □

The equilibria in (46), feature two sources of inefficiency. First, aggregate output falls short of the efficient level A_0. Second, since price estimates vary, output and the marginal cost of output differ across firms. Aggregate output is thus produced inefficiently.

References

Ackerman, N., Freer, C., and Roy, D. (2011). Noncomputable conditional distributions. *Proceedings - Symposium on Logic and Computer Science, Toronto*, pages 31–36.

Arrow, K. J. (1970). *Essays in the Theory of Risk Bearing*. Amsterdam: North-Holland Publishing.

Arthur, B. (2015). *Complexity and the Economy*. Oxford University Press.

Blanchard, O. J. (1979). Backward and forward solutions for economies with rtional expectations. *American Economic Review Papers and Proceedings*, 69:114–118.

Bray, M. (1982). Learning, estimation, and stability of rational expectations. *Journal of Economic Theory*, 26:318–339.

Chamley, C. (2014). When demand creats its own supply: savings traps. *Review of Economic Studies*, 81(2):651–680.

Chiang, A. C. and Wainwright, K. (2005). *Fundamental Methods of Mathematical Economics*. Boston [u.a.]: McGraw-Hill, 4. internat. edition.

Cooper, R. and John, A. (1988). Coordinating coordination failures in Keynesian models. *Quarterly Journal of Economics*, 103(4):441–463.

DeCanio, S. (1979). Rational expectations and learning from experience. *Quarterly Journal of Economics*, 93:47–57.

Diamond, P. A. (1965). National debt in a neoclassical growth model. *American Economic Review*, 55(5):1126–1150.

Diamond, P. A. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy*, 90(5):881–894.

Farmer, D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature, Opinion*, 460(6):685–686.

Finetti, B. D. (1952). Sulla preferibilita. *Annali di Economica*, (11):685–709.

Galor, O. (2007). *Discrete Dynamical Systems*. Springer Verlag.

Galton, F. (1907). Vox populi. *Nature*, 75:450–451.

Gintis, H. (2007). The dynamics of general equilibrium. *Economic Journal*, 117:1280–1309.

Grossman, S. and Stiglitz, J. (1976). Information and competitive price systems. *American Economic Review Papers and Proceedings*, 66(2):246–253.

Grunberg, E. and Modigliani, F. (1954). The predictability of social events. *Journal of Political Economy*, 62:465–478.

Hellwig, M. (1993). The conceptual structure of macroeconomic models: I the income equation. *WWZ Working paper 9308 University Basel*.

Hofbauer, J. and Sandholm, W. (2011). Survival of dominated strategies under evolutionary dynamics. *Theoretical Economics*, 6(3):341–377.

Hutchison, T. (1937). Expectation and rational conduct. *Journal of Economics*, 8(3):636–653.

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. New York: Harcourt-Brace Inc.

Kuhle, W. (2016). Darwinian adverse selection. *Algorithmic Finance*, 5:31–36.

Marcet, A. and Sargent, T. (1989). Convergence of least squares learning mechanisms in self-refrential linear stochastic models. *Journal of Economic Theory*, 48:337–368.

Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.

Maymin, P. (2011). Markets are efficient if and only if p=np. *Algorithmic Finance*, 1(1):1–11.

McLennan, A. (1984). Price dispersion and incomplete learning in the long run. *Journal of Economic Dynamics and Control*, 7:331–347.

Morris, S. and Shin, H. S. (1998). Unique equilibrium in a model of self-fullfilling currency attacks. *American Economic Review*, 88(3):578–597.

Muth, J. (1961). Rational expectations and the theory of price movements. *Econometrica*, 29:315–335.

Nelson, R. and Winter, S. (1985). *An Evolutionary Theory of Economic Change*. Harvard University Press.

Niederhoffer, V. (1997). *The Education of a Speculator*. John Wiley.

Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1):122–136.

Rothschild, M. (1974). A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9:185–202.

Rubinstein, A. (1998). *Modeling Bounded Rationality*. MIT press.

Samuelson, P. A. (1947). *Foundations of Economic Analysis*. Cambridge Ma.: Harvard University Press.

Samuelson, P. A. (1962). Parable and realism in capital theory: The surrogate production function. *Review of Economic Studies*, (29):193–206.

Samuelson, P. A. (2009). *Economics: An Introductory Analysis*. McGraw-Hill, 19th. edition.

Sargent, T. (1993). *Bounded Rationality in Macroeconomics*. Oxford University Press.

Soros, G. (1994). *The Alchemy of Finance*. John Wiley.

Tesfatsion, L. (2006). Agent based computational economics: A constructive approach to economic theory. In *Handbook of Computational Economics*, pages 831–880. ed. L. Tesfatsion, K.Judd, Elsevier.

Thurner, S., Farmer, D., and Geanakoplos, J. (2012). Leverage causes fat tails and clustered volatility. *Quantitative Finance*, 12(5):695–707.

Ulanowicz, R. (2008). The dual nature of ecosystem dynamics. *Ecological Modelling*, 220:1886–1892.